UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/643,585	08/18/2003	Steven L. Scott	1376.700US1	4004

21186 7590 01/25/2007
SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
P.O. BOX 2938
MINNEAPOLIS, MN 55402

EXAMINER	
TSAI, SHENG JEN	
ART UNIT	**PAPER NUMBER**
2186	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	01/25/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary

Application No.	Applicant(s)
10/643,585	SCOTT, STEVEN L.
Examiner	**Art Unit**
Sheng-Jen Tsai	2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *04 December 2006*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1,3-8 and 11-18* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1,3-8 and 11-18* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *18 August 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)

2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *12/04/2006*.

4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

5) ☐ Notice of Informal Patent Application

6) ☐ Other: _____.

DETAILED ACTION

1. This Office Action is taken in response to Applicants' Amendments and Remarks filed on 12/04/2006 regarding Application 10,643,585 filed on 08/18/2003.

2. Claims 2 and 9-10 have been cancelled previously.

Claims 1, 8, 12, 15 and 17 have been amended.

Claims 1, 3-8 and 11-18 are pending under consideration.

3. ***Response to Remarks and Amendments***

Applicants' amendments and remarks have been fully and carefully considered, with Examiner's responses set forth below.

<u>Amendments on Claims 1, 8, 12 and 17</u>

Each of the independent claims 1, 8, 12 and 17 has been amended with additional limitations of "a processor cache and a translation look-aside buffer (TLB)," "the RTT has capacity to store all physical page numbers associated with the processing node," and "wherein each TLB translates memory references from its associated processor to the shared memory within the processing node."

However, the Specification and Figures of Applicant's disclosure are completely silent regarding the element of "translation look-aside buffer (TLB)," and Examiner was not able to locate or identify any citing or description of TLB. As such, the newly added limitations of "a translation look-aside buffer (TLB)" and "wherein each TLB translates memory references from its associated processor to the shared memory within the processing node" lack the support by the written description.

Similarly, the Specification and Figures of Applicant's disclosure are completely silent regarding the element of "physical page numbers," and Examiner was not able to locate or identify any citing or description of "physical page numbers." As such, the newly added limitation of "the RTT has capacity to store all physical page numbers associated with the processing node" lacks the support by the written description.

Remarks on Differences between Application and Reference

Applicant contends that the address translation scheme for handling a virtual address destined for a local node and a remote node is different between the Application and the cited reference (Scott et al., US 6,925,547, hereafter referred to as Scott). The examiner disagrees with this assessment due to the following reasons:

First, Applicant's Specification (page 6, lines 19-30 and page 7, lines 1-8, filed on 08/18/2003) states that

"Figure 10 shows one embodiment of local memory 1000 used in the processing node 500 of Figure 5. In this embodiment, local memory includes two MSP ports 1010, two Cache Coherence Directories 1040, a crossbar switch 1020, two network ports 1030, a Remote Address Translation Table (RTT) 1050, and RAM 1060. Remote Translation Table (RTT) 1050 translates addresses originating at remote processing nodes 500, 600, 700, 800 to physical addresses at the local node. In some embodiments, this includes providing a virtual memory address at a source node, determining that the virtual memory address is to be sent to a remote node, sending the virtual memory address to the remote node, and translating the virtual memory address on the remote node into a physical memory address using a RTT. The RTT contains

translation information for an entire virtual memory address space associated with the remote node. Another embodiment of RTT provides for translating a virtual memory address in a multi-node system. The method includes providing a virtual memory address on a local node by using a virtual address of a load or a store instruction, identifying a virtual node associated with the virtual memory address, and determining if the virtual node corresponds to the local node. If the virtual node corresponds to the local node, then the method includes translating the virtual memory address into a local physical memory address on the local node. If, instead, the virtual node corresponds to a remote node, then the method includes sending the virtual memory address to the remote node, and translating the virtual memory address into a physical memory address on the remote node."

From the description, it appears that the address translation scheme would use the local translation means (i.e., the RTT located at the local processor/node) to translate the address if the virtual address corresponds to a local address space, and would use the remote translation means (i.e., the RTT located at the remote processor/node) to translate the address if the virtual address corresponds to a remote address space.

Second, the reference teaches the remote address translation as follows "As described above, the local TLB can be used by a local CE 64 to perform translations for local memory accesses, thereby allowing the user to program the CE using virtual addresses. As now described, CE 64 can also be programmed to send virtual addresses to a remote or target node for remote memory accesses (using the

CD associated with the virtual address to identify the remote node), with the TLB on that node being used to translate those addresses (column 17, lines 35-45; see also column 2, lines 65-67 and column 3, lines 1-23)."

Thus, it appears that the address translation scheme would use the local translation means (i.e., the RTT located at the local processor/node) to translate the address if the virtual address corresponds to a local address space, and would use the remote translation means (i.e., the RTT located at the remote processor/node) to translate the address if the virtual address corresponds to a remote address space.

Therefore, the Examiner does not see the difference suggested by Applicant's remarks.

Claim Rejections - 35 USC § 112

4. The following is a quotation of the first paragraph of 35 U.S.C. 112:

> The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

5. Claims 1, 3-8 and 11-18 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

Each of the independent claims 1, 8, 12 and 17 has been amended with additional limitations of "a processor cache and a translation look-aside buffer (TLB)," "the RTT has capacity to store all physical page numbers associated with the

processing node," and "wherein each TLB translates memory references from its associated processor to the shared memory within the processing node."

However, the Specification and Figures of Applicant's disclosure are completely silent regarding the element of "translation look-aside buffer (TLB)," and Examiner was not able to locate or identify any citing or description of TLB. As such, the newly added limitations of "a translation look-aside buffer (TLB)" and "wherein each TLB translates memory references from its associated processor to the shared memory within the processing node" lack the support by the written description.

Similarly, the Specification and Figures of Applicant's disclosure are completely silent regarding the element of "physical page numbers," and Examiner was not able to locate or identify any citing or description of "physical page numbers." As such, the newly added limitation of "the RTT has capacity to store all physical page numbers associated with the processing node" lacks the support by the written description.

Claims 3-7, 11, 13-16 and 18 are rejected by virtue of their dependency from claims 1, 8, 12 and 17, respectively.

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

7. Claims 1, 3 and 5-8 and 11-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Scott et al. (US 6,925,547), and in view of Fossum et al. (US 4,888,679).

As to claim 1, Scott et al. disclose **a computer system** [figures 1-3] **comprising**:

a network [interconnection network, figure 2, 14],

one or more processing nodes connected via the network [figures 1-3], **wherein each processing node includes**:

a plurality of processors [PM, figure 1, 12], **wherein each processor includes a scalar processing unit, a vector processing unit and means for operating the scalar processing unit independently of the vector processing unit** [taught by Fossum et al., see below],

a processor cache [column 5, lines 48-53; To support local address translations, each SHUB contains a translation-lookaside buffer (TLB) 108 for performing local address translations for both block transfers and AMOs. <u>A TLB is a cache</u> that holds only page table mappings (column 16, lines 7-15)] **and a translation look aside buffer (TLB)** [abstract; column 1, lines 40-53; To support local address translations, each SHUB contains a translation-lookaside buffer (TLB) 108 for performing local address translations for both block transfers and AMOs. A TLB is a cache that holds only page table mappings (column 16, lines 7-15)], **wherein the scalar processing unit places instructions for the vector processing unit in a queue foe execution by the vector processing unit and the scalar processing unit continues to execute additional instructions** [taught by Fossum et al., see below]; **and**

a shared memory connected to each of the processors within the processing node [memory, figure 2, 28A and 28B; column 5, lines 47-67], **wherein the shared memory includes a cache** [To support local address translations, each SHUB contains a translation-lookaside buffer (TLB) 108 for performing local address translations for both block transfers and AMOs. A TLB is a cache that holds only page table mappings (column 16, lines 7-15)] **and a Remote Address Translation table (RTT), wherein the RTT has capacity to store all physical page numbers associated with the processing node** [column 17, lines 35-45; column 2, lines 65-67 and column 3, lines 1-23] **and wherein the RTT translates memory addresses received from other processing node such that the memory addresses are translated into physical addresses within the shared memory** [A method of performing remote address translation in a multiprocessor system includes determining a connection descriptor and a virtual address at a local node, accessing a local connection table at the local node using the connection descriptor to produce a system node identifier for a remote node and a remote address space number, communicating the virtual address and remote address space number to the remote node, and translating the virtual address to a physical address at the remote node (qualified by the remote address space number) (abstract); figures 4A, 4B, 5A and 5B; column 25, lines 39-50];

wherein processors on one node can load data directly from and store data directly to shared memory on another processing node via addresses that are

translated on the other processing node using the other processing node's RTT [abstract; figures 4A, 4B, 5A and 5B; column 25, lines 39-50]; **and**

wherein each TLB translates memory references from its associated processor to the shared memory within the processing node [As described above, the local TLB can be used by a local CE 64 to perform translations for local memory accesses, thereby allowing the user to program the CE using virtual addresses. As now described, CE 64 can also be programmed to send virtual addresses to a remote or target node for remote memory accesses (using the CD associated with the virtual address to identify the remote node), with the TLB on that node being used to translate those addresses (column 17, lines 35-45; see also column 2, lines 65-67 and column 3, lines 1-23)].

Regarding claim 1, Scott et al. do not teach that **each processor includes a scalar processing unit, a vector processing unit and means for operating the scalar processing unit independently of the vector processing unit**.

However, the concepts of scalar processors and vector processors is well known and widely used in the art. Essentially every PC has a scalar processor for data processing, and vector processors are commonly used for graphic applications (see Microsoft Computer Dictionary, 5th edition, 2002, Microsoft Press, page 548 – vector and page 549 – vector graphics).

Further, Fossum et al. disclose in their invention "Method and Apparatus Using a Cache and Main memory for Both Vector Processing and Scalar Processing by Prefetching Cache Blocks Including Vector Data Elements" an apparatus comprising a

vector processor (figure 1, 22; figure 7, 116) and a scalar processor (figure 1, 21; figure 7, 108) where the scalar processor and the vector processor operate independently of each other (figure 7; column 2, lines 35-68; column 3, lines 1-43). Including both scalar and vector processors in a computer system with a cache allows the prefetching of block data using the vector processor and increases the data throughput (column 2, lines 12-34).

Specifically, Fossum et al. disclose that **each processor includes a scalar processing unit, a vector processing unit and means for operating the scalar processing unit independently of the vector processing unit** [a vector processor (figure 1, 22) is added to a digital computing system 9figure 1, 20) including a scalar processor (figure 1, 21), a virtual address translation buffer, a main memory (figure 1, 23), and a cache (figure 1, 24) (column 3, lines 7-10); figure 7 shows the detailed organization of these components], **wherein the scalar processing unit places instructions for the vector processing unit in a queue for execution by the vector processing unit** [Another object of the invention is to take a main memory and cache optimized for scalar processing and make it suitable for vector processing as well (column 2, lines 40-42); in accordance with the invention, a main memory and cache suitable for scalar processing are used in connection with a vector processor by issuing prefetch requests in response to the recognition of a vector load instruction (column 2, lines 47-51); In response to a vector load instruction, <u>the scalar processor executes microcode for sending a vector load command to the vector processor</u>, and also for sending the vector prefetch requests to the cache. The vector prefetch requests

include the virtual addresses of the blocks that will be accessed by the vector processor. These virtual addresses are computed based upon the vector address, the length of the vector, and the stride or spacing between the addresses of the adjacent elements of the vector (column 3, lines 17-26); FIG. 7 is a preferred embodiment of the present invention which uses microcode in a scalar processing unit to generate vector prefetch requests for an associated vector processing unit (column 3, lines 67-68); column 11, lines 35-46] **and the scalar processing unit continues to execute additional instructions** [Specifically, the scalar processing unit includes a micro-sequencer and issue logic 109 which executes prestored microcode 110 to interpret and execute the parsed instructions from the instruction processing unit 107. These instructions include <u>scalar instructions</u> which the micro-sequencer and issue logic executes by operating a register file and an arithmetic logic unit 111. These scalar instructions include, for example, an instruction to fetch scalar data from the cache unit 106 and load the data in the register file 111 (column 11, lines 35-46)].

It is well known in the art that the use of vector processors increases the throughput by processing multiple vector elements simultaneously as opposed to processing a single element at a time.

Therefore, it would have been obvious for one of ordinary skills in the art at the time of Applicant's invention to recognize the benefit of having both scalar and vector processing units, as demonstrated by Fossum et al., and to incorporate it into the existing apparatus disclosed by Scott et al. to further enhance the performance of the system.

As to claim 3, Scott et al. teach that **the shared memory further includes a plurality of cache coherence directories, wherein each processing node is coupled to one of the cache coherence directories** [In one embodiment, all of the coherence information is passed across the bus in the form of messages, and each processor on the bus "snoops" by monitoring the addresses on the bus and, if it finds the address of data within its own cache, invalidating that cache entry. Other cache coherence schemes can be used as well (column 5, lines 47-67)].

As to claim 5, Scott et al. teach that **the processing nodes include at least one input/out (I/O) channel controller** [I/O, figure 1, 18], **wherein each I/O channel controller is coupled to the shared memory of the processing node** [figures 1-3; column 4, lines 10-22].

As to claim 6, Fossum et al. teach that **each scalar processing unit contains a scalar cache memory** [cache, figure 1, 24 is associated and shared by the scalar (21) and vector (22) processing units], **wherein scalar cache memory contains a subset of cache lines stored in the shared memory cache** [column 4, lines 15-54];

a plurality of address latches each of which for outputting register set address bit by latching a address, in response to the register set control signal and the self-refresh signal when the mode register set signal is applied [column 8, lines 3-18]; **and**

a partial array self-refresh controller for selectively activating the plurality of control signals by decoding the plurality of register set addresses depending on

input of the internal address [the refresh controller, figure 2, 217; column 6, lines 39-45].

As to claim 7, Scott et al. teach that **the network includes a router connecting one or more of the processing nodes** [R (Router), figure 1, 16] .

As to claim 8, refer to "As to claim 1."

As to claim 11, refer to "As to claim 3."

As to claim 12, refer to "As to claim 1."

As to claim 13, refer to "As to claim 3."

As to claim 14, refer to "As to claim 5."

As to claim 15, refer to "As to claim 6."

As to claim 16, refer to "As to claim 7."

As to claim 17, refer to "As to claim 1."

As to claim 18, refer to "As to claim 3."

8. Claim 4 is rejected under 35 U.S.C. 103(a) as being unpatentable over Scott et al. (US 6,925,547), in view of Fossum et al. (US 4,888,679), and further in view of Nakazato (US 6,782,468).

As to claim 4, neither Scott et al. nor Fossum et al. teach that **each processor includes two vector pipelines**. However, Nakazato discloses in the invention "Shared Memory Type Vector Processing Syatem, Including a Bus for Transferring a Vector Processing Instruction, and Control Method Thereof" an apparatus comprising multiple vector pipelines in each processor (n vector processing units, figure 2, 14a~14n) and a scalar processor (figure 2, 11). Including multiple vector processors in a computer

system allows the multiple vector processing tasks to be performed simultaneously and increases the data throughput. Therefore, it would have been obvious for one of ordinary skills in the art at the time of Applicant's invention to recognize the benefit of having multiple vector processing units, as demonstrated by Nakazato, and to incorporate it into the existing apparatus disclosed by Scott et al. and Fossum et al. to further enhance the performance of the system.

9. ***Related Prior Art***

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Schimmel, (US 6,105,113), "System and Method for Maintaining Translation Look-Aside Buffer (TLB) Consistency."
- Scott, (US 6,922,766), "Remote Translation Mechanism for a Multi-Node System."
- Nesheim et al., (US 5,897,664), "Multiprocessor System Having Mapping Table in Each Node to Map Global Physical Addresses to Local Physical Addresses of Page Copies."
- Vishin et al., (US 5,860,146), "Auxiliary Translation Lookaside Buffer for Assisting in Accessing Data in Remote Address Space."
- Deneau, (US 6,684,305), "Multiprocessor System Implementing Virtual Memory Using a Shared Memory, and a Page Replacement Method for Maintaining Paged memory Coherence."

- Frank et al., (US 6,490,671), "System for Efficiently Maintaining Translation Lookaside Buffer Consistency in a Multi-Threaded, Multi-Processor Virtual Memory System."
- Hansen, (US 6,101,590), "Virtual Memory System with Local and Global Virtual Address Translation."

Conclusion

10. Claims 1, 3-8 and 11-18 are rejected as explained above.

11. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

10. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

January 18, 2007

[illegible]
[illegible]
1/21/07